A Prototype Web Site for Immunization Knowledge Maintenance

Perry L. Miller, M.D., Ph.D., Sandra J. Frawley, Ph.D.
Cynthia Brandt, M.D., M.P.H., Frederick G. Sayward, Ph.D.
Center for Medical Informatics
Yale University School of Medicine
New Haven, CT

IMM/Web is a prototype Web site designed to assist in the maintenance of the knowledge required to perform computer-based forecasting for childhood immunization. IMM/Web operates in conjunction with IMM/Serve, a immunization forecasting program that takes a child's immunization history and produces recommendations as to which vaccinations are due, and which should be scheduled next. IMM/Serve's domain knowledge is expressed in both tabular and rule-based form. Using IMM/Web, the various tabular forecasting parameters can be modified via the Web. Test cases can then be generated automatically which can be used to help verify the new version of logic. Finally, the test cases can be automatically passed for IMM/Serve to analyze using the newly defined parameters. The IMM/Web project is exploring how the process of updating, customizing, and testing new versions of a computer-based clinical guideline might be performed and guided in an organized fashion via the World-Wide Web.

INTRODUCTION

IMM/Web is a prototype Web site designed to assist in the process of modifying and testing different versions of a computer-based guideline for childhood immunization: (http://ycmi.med.yale.edu/immserve/). The clinical knowledge base required to produce patient-specific immunization recommendations is surprisingly complex. The recommendations themselves change quite frequently as the field evolves. In addition, individual immunization registries may wish to use somewhat customized versions of the logic. As a result, a major challenge for the field will be to maintain different versions of the knowledge in computer-based form, and to assure that the computer-based guideline operates correctly with all versions of the knowledge.

IMM/Web has been built to explore how the World-Wide Web might help deal with this problem. As described previously [1], we have built IMM/Serve, a program which performs immunization forecasting. IMM/Serve takes as input a child's immunization history (which may be extracted from an immunization registry database [2]) and produces recommendations as to which immunizations are due as of a specified date, and which should be scheduled next.

We are currently working with several organizations, including Yale New Haven Medical Center, the State of Oregon, and the US Indian Health Service to explore the issues involved in maintaining this knowledge. It is clear that different organizations will typically want different versions of the logic, and may even want different versions running in different clinics. As described later in more detail, IMM/Web is designed to help maintain different versions of the immunization knowledge in several ways.

1. IMM/Web first allows the user to perform Web-based modification of immunization forecasting parameter tables, which include the minimum ages and wait-intervals for the various doses of each vaccine series.

2. IMM/Web then processes IMM/Serve's rule-based knowledge together with the new parameters to generate a set of test cases designed to help test the newly defined version of the knowledge.

3. Finally, IMM/Web automatically passes the test cases to IMM/Serve for analysis using the newly defined parameters. This helps the user verify that the new system as a whole performs correctly.

4. The user will ultimately also be able to download the new, customized version of the tables to link to IMM/Serve, for example, to run as a component of a local immunization registry system.

Thus, IMM/Web is designed to serve several roles. It is designed to help coordinate the dissemination of different versions of an immunization forecasting program. It is also designed to provide a centrally

1091—8280/97/$5.00 © 1997 AMIA, Inc.

maintained resource to help test and validate different versions of the knowledge.

BACKGROUND

It is a major task to build and test the initial computer-based version of a complex clinical guideline designed to produce patient-specific recommendations. It is widely acknowledged, however, that the task of maintaining such a knowledge-based system over time is at least equally challenging [3,4]. A variety of projects have explored these issues, at least in part, often in the context of making the knowledge reusable or sharable, or in the context of developing tools to acquire the knowledge automatically [5-7]. One problem is that the clinical field itself evolves so that the recommendations of national guideline panels change. As a result, not only must the program be modified, but if the modifications are significant, a major process of reverification may be required. These problems are compounded if the users of the program need to customize the logic to meet their particular needs.

The problem of guideline knowledge maintenance may well turn out to be the biggest challenge in the development of interactive, computer-based versions of complex clinical guidelines. The IMM/Web project is exploring how the Web might be used as a vehicle to help meet these challenges.

THE IMM/WEB BUILDING BLOCKS: IMM/SERVE, IMM/DEF, AND IMM/TEST

IMM/Serve

IMM/Serve is an immunization forecasting program that contains the guidelines for 7 vaccination series commonly administered to children [8,9]: Diphtheria Tetanus Pertussis, Haemophilus influenzae type b, Hepatitis A, Hepatitis B, Measles Mumps Rubella, Polio, and Varicella. IMM/Serve stores its domain knowledge in two general formats.

1. Tabular knowledge Tables are used to store the various parameters used for immunization forecasting: for example, the minimum, recommended, and "past due" ages for each dose in a vaccination series, and the minimum and "past due" wait-intervals between doses. This information is stored as tabular parameter sets. For a given dose of a vaccine, there may be several different parameter sets, each for use in different clinical conditions. IMM/Serve currently contains roughly 50 different forecasting parameter sets for the seven vaccination series.

   In addition, tabular knowledge in the form of a decision table is used to represent the screening logic that determines whether a previous dose has been given too soon to be counted (e.g., at too early an age or too soon after a previous dose).

2. Rule-based knowledge If-then rules are used to represent the clinical knowledge that determines which set of forecasting parameters applies in a particular case. For certain vaccine series, rules are also used to determine the vaccine brand or preparation which is indicated. There are currently over 250 if-then rules in IMM/Serve.

   IMM/Serve's rule-base contains different versions of certain vaccine recommendations. (A flag passed to IMM/Serve indicates which version of the rules to use in processing a case.) These different versions reflect recent changes to the recommendations made by the national panel. A new version of the guideline will typically be adopted by different clinics at different times (for example, because of different policies and regulations in different States and in different health organizations). As a result, at any given time, different clinics may be using different versions of the rules.

IMM/Serve is built using the C programming language. Its if-then rules are expressed using KnowledgeCraft, a rule-based shell program which translates the rules so that they also can be executed by a C program.

IMM/Def

The rule-based knowledge contained in IMM/Serve is quite complex in part because it must respond to all possible combinations of patient presentation for each dose, and must react to these in three different contexts: a) when the dose is "due now", b) when the dose is "not yet due", and c) when the dose must be "scheduled to follow" a dose that is due now.

To allow IMM/Serve to react to these three contexts for each dose, the same underlying logic (which we call "immunization definition logic") must be adapted to fit the three contexts. As a result, the if-then rules can be thought of as "implementation logic" which express the underlying definition logic in different

ways to accommodate the three contexts in which it must be applied.

To facilitate the creation and validation of IMM/Serve's rule-based implementation logic, we have created IMM/Def, a Lisp program which takes as input the relatively simply expressed definition logic for a vaccine series and automatically generates the more complex rule-based implementation logic for the most complex portion ("kernel") of the rules. In this translation process, certain variable names must be changed to reflect the different temporal contexts in which the logic must operate. IMM/Def is currently used as a tool to help double check the "kernel" portion of IMM/Serve's rules, but could in the future be used to automatically generate those rules from scratch.

IMM/Test

The immunization definition logic developed for IMM/Def is also used by IMM/Test, a Lisp program built to generate a set of test cases. These test cases are designed to help assure that the system as a whole operates correctly for the different combinations of conditions which might arise. To generate test cases, IMM/Test goes through several steps for each dose of a vaccine series.

1. IMM/Test processes the definition logic for the dose to determine the sets of conditions to which the logic must respond.

2. It then generates all possible combinations of those conditions. In the process it removes any combinations of conditions that do not make sense, for example, a combination of conditions that included both "age < 4 months" and "age >= 17 years".

3. For each meaningful combination of conditions, IMM/Test processes the definition logic to determine what parameter set is indicated for that combination. If for some reason, no parameter set is indicated, the combination is flagged for manual inspection to make sure that it does not represent a "hole" in the logic.

4. For each meaningful combination of conditions, IMM/Test then generates one or more test cases that meet all those conditions when processed by IMM/Serve. The construction of a test case frequently requires a certain amount of domain knowledge to allow IMM/Test, for example, to include a set of previous vaccination dates that will appropriately activate the conditions identified. For each combination of conditions, IMM/Serve produces 1) a case that satisfies the forecasting parameter set, and usually 2) one or more cases which fail to satisfy the forecasting parameter set for different reasons. For example, the forecast date might be too early for the dose based on the child's age, or too early based on the wait-interval from the previous dose.

In this way IMM/Test produces a set of test cases that is "comprehensive" in the sense that, for each dose in a vaccine series, it activates the different combinations of variables expressed in the "kernel" of IMM/Serve's rule-based logic and in the associated tabular forecasting parameters. The test cases are produced in a format that can be directly input to IMM/Serve for analysis.

IMM/WEB IN OPERATION

IMM/Web is implemented in its current prototype form to demonstrate certain basic capabilities that we believe will assist in immunization knowledge maintenance. The approach will be refined and extended as we gain experience with its use. IMM/Web builds directly on the three tools described in the previous section (IMM/Serve, IMM/Def, and IMM/Test). In this section, we describe the current operation of IMM/Web. (See Figure 1.)

1. To use IMM/Web, the user goes to a JavaScript-based Web page that presents IMM/Serve's forecasting parameters for a given vaccine in tabular form. These indicate the minimum ages and wait-intervals that determine IMM/Serve's operation. The user may edit these tables, making changes to the parameters. The user is not currently able to change the underlying rule-based logic which determines the clinical conditions in which the various parameter sets will be used. The user may, however, select among different versions of the rule-based logic which are already built into IMM/Serve.

2. Once the user has made a set of changes to the forecasting parameters, he can then request (by clicking on a button) that the parameters be passed to IMM/Test for automatic generation of test cases. If a choice exists, the user specifies which version of the rules he wishes to use in the analysis.

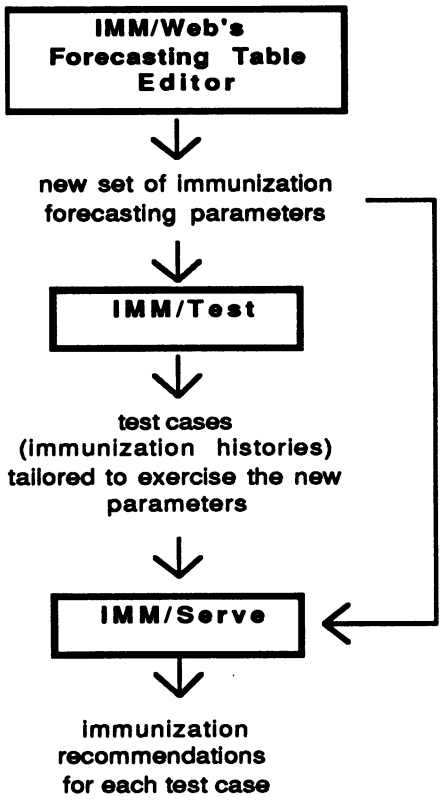

Figure 1: Overview of IMM/Web's operation in its current prototype form.

3. The new parameters are passed to IMM/Test which produces a tailored set of test cases as described above. The test cases are displayed to the user, who may download them or alternatively may ask (by clicking a button) that they be automatically processed by IMM/Serve directly on the Web to produce patient-specific immunization recommendations.

4 If automatic processing by IMM/Serve is requested, the newly defined parameter table is passed to the IMM/Serve program and the test cases are processed. Web output is produced showing each case and IMM/Serve's recommendations for that case.

The purpose is to help the user verify that the system as a whole is operating correctly with the newly defined parameters. In addition, IMM/Web can be used to validate the operation of any other immunization forecasting program, whose analysis on the test cases can be compared to IMM/Serve's. We emphasize that IMM/Web is merely a tool designed to assist in the process of immunization knowledge maintenance. There are many aspects of a program like IMM/Serve that need to be tested. IMM/Web currently demonstrates one approach to helping perform certain aspects of this task.

FUTURE DIRECTIONS

We anticipate that the current IMM/Web prototype will be extended in the future in a number of ways.

1. It will be important to link IMM/Web to a database. Then, as users customize the knowledge, these changes can be stored at the Web site in an organized, password-protected fashion.

2. In addition to allowing the user to modify the forecasting parameters, it will also be important to permit the user to modify other parameters, including those used for screening each dose in a child's immunization history to determine if it was given too early to be counted.

3. It may also be useful to allow IMM/Web users to create their own versions of the if-then rules, and to use the IMM/Def program to help double check their new rule kernel. This capability would involve a major programming effort to implement. In addition, using the capability would require a very computer-sophisticated user since creating and validating a comprehensive set of rules for a vaccine series is a difficult task.

4. There are a variety of other aspects of IMM/Serve's functionality whose testing could also be assisted by automatically generated test cases. These include a) testing the handling of potential interactions between different live vaccines, b) assuring that the integration of wait intervals, ages, and dates is performed correctly, and c) testing aspects of the rule-based logic not included in the kernel rules.

5. Tools could be provided that automatically scan the test case forecasts looking for certain types of errors, to help the user in this analysis.

6. A further important step will involve refinement of the system in the context of practical operational use to see how it can best meet the needs of real users.

DISCUSSION

One of the major problems that arises in the widespread dissemination of complex clinical software which needs to be updated and customized, is how to control the correctness and currency of different versions of the system. For example, if the software is disseminated via diskettes, or downloaded over the Internet, and tools are provided to modify the knowledge locally by users on their own machines, it will be difficult to assure that all versions of the software are updated appropriately.

Although IMM/Serve itself will typically need to run locally as part of a user's immunization registry, the IMM/Web approach allows the tools used to maintain and test the knowledge to be kept centrally in one place (at the Web site).

1. These tools can be accessed readily from any machine connected to the Internet which possesses a Web browser.

2. The Web site provides a vehicle to make new versions of the computer-based knowledge rapidly available to immunization registries.

3. The Web also provides an easy mechanism to make an evolving suite of immunization knowledge maintenance tools available.

4. As minor changes are made to the knowledge or to the tools, those changes can be readily disseminated if they are always used via the Web.

5. As registries create their own customized versions of the knowledge, that knowledge can be stored at the Web site. Then if any changes are made to the underlying guidelines, the registries could be automatically notified that their customized versions may require change.

We feel that the Web represents an important technology that will significantly facilitate the type of knowledge maintenance which we are proposing. The current IMM/Web prototype is a first step in exploring how this capability might best be accomplished.

Acknowledgments

This work was supported in part by NIH grants T15 LM07056 and G08 LM05583 from the National Library of Medicine.

References

1. Miller PL, Frawley SJ, Sayward FG, Yasnoff WA, Duncan L, Fleming D: IMM/Serve: A rule-based program for childhood immunization. Proceedings of the Annual Fall Symposium of the American Medical Informatics Association (Formerly SCAMC), Washington DC, October 1996, pp. 184-188.
2. Gostin LO, Lazzarine Z. Childhood immunization registries. JAMA, 1995; 274:1793-1799.
3 Coenen F, Bench-Capon T. Maintenance of Knowledge-Based Systems. Academic Press, New York, 1993.
4. Shiffman RN. Guideline maintenance and revision: Fifty years of Jones Criteria for the diagnosis of rheumatic fever. Archives of Pediatrics and Adolescent Medicine, 1995; 149:727-732.
5. Shahar Y, Miksch S, Johnson P. An intention-based language for representing clinical guidelines. Proceedings of the Annual Fall Symposium of the American Medical Informatics Association (Formerly SCAMC), Washington DC, October 1996, pp. 592-596.
6. Hripcsak G, Ludemann P, Pryor TA, Wigertz OB, Clayton PD. Rationale for the Arden Syntax. Computers and Biomedical Research, 1994, 27:291-324
7. Musen MA, Gennari JH, Eriksson H, Tu SW, Puerta AR. PROTÉGÉ-II: Computer support for development of intelligent systems from libraries of components. In: RA Greenes, HE Peterson, DJ Protti, (Eds.). Proceedings of Medinfo'95. North Holland, New York, 1995, pp. 766-770.
8. Centers for Disease Control and Prevention: Recommended Childhood Immunization Schedule - United States, January - June 1996. MMWR 44 (Nos. 51 & 52):942-943, 1996.
9. Atkinson W, Gandtt J, Mayfield M, Furphy L, (Eds.). Epidemiology and Prevention of Vaccine-Preventable Diseases. CDC, Atlanta, GA, 1995.